

Replies to Commentators

Keith DeRose

Yale University

keith.derose@yale.edu

Abstract

Replies are given to comments, questions, and objections to *The Appearance of Ignorance*. The reply to Robin McKenna focuses mainly on his questions of whether, with the skeptical argument I'm focused, a strong enough appearance of ignorance is generated to require an account of that appearance, and whether, to the extent that we do need to account for that appearance, we might do so without contextualism by adopting a solution proposed by Ernest Sosa. The reply to Michael Blome-Tillman focuses mainly on a counterexample he offers to my account of the operation of the "Rule of Sensitivity." The reply to Elke Brendel focuses mainly on objections to contextualism from the phenomena of disagreement and retraction. The reply to Peter Baumann concerns several of his comments about my treatment of the Harman lottery puzzle.

Keywords

Skepticism, counterexamples, epistemic conservatism, lotteries, safety, sensitivity

1. Reply to McKenna: On the Sosa Safety Solution

As Robin McKenna nicely summarizes his critical aims, they are to “argue *inter alia* that we don’t need a contextualist semantics to account for the appearance of ignorance, and in any case that the ‘strength’ of the appearance of ignorance is unclear, as is the need for a philosophical diagnosis of it.” I will reply to these two arguments in reverse order.

McKenna’s criticism of the power of the appearance of ignorance generated by the skeptical argument I’m focused on, AI, breaks the matter down into the issues of how widespread, robust, and rational that appearance is. He isn’t so much arguing in his own voice that the power isn’t significant, but using features of my own conclusions about that matter against me. The flavor of his approach is conveyed well in this preview of his case:

As we will see, somewhat surprisingly, DeRose doesn’t think that the appearance of ignorance is particularly widespread, robust, or rational. In this section, I will outline why DeRose thinks this, and explain why I take this to pose a problem for the central argument of [*The Appearance of Ignorance*]. The problem, in short, is that if the appearance of ignorance isn’t that widespread, robust, or rational, why should a solution to AI have to account for it?

And he has a good deal of ammunition for the making of the case, because, when I assess the power of the skeptical argument in sections 2.7-2.11 of the book (at 2017: 53-63), I work hard to give a very balanced account. I should add that while I think I achieve a balanced and responsible assessment, I don’t take it to be anything like a definitive, final one, but one that should have the flavor of one based on a preliminary inquiry. I think philosophy as a discipline has a long way to go in learning how to address such questions (about such things as the intuitive power of claims or arguments), and, in particular, how to get and to use empirical results in doing so. But I do think

that even a preliminary look provides very strong grounds for the fairly minimal, and, I think, quite safe, conclusion that the power of AI, and, particular, its first premise is such as to be worth accounting for.

Where O is something we would ordinarily take ourselves to know (e.g., *I have hands*), and H is a suitably chosen skeptical hypothesis (e.g., *I am a BIV*), our skeptical argument is:

The Argument from Ignorance (AI):

1. I don't know that not-H
2. If I don't know that not-H, then I don't know that O

So, C. I don't know that O

My assessment of the power of the argument takes place against the background of what seems to be the strong sense among current epistemologists that premise 1 is the argument's weak link, and so is focused on the intuitive power of that premise.

As I report in section 2.7 of the book, there has been a sharp division among epistemologists about the intuitive power of that first premise, with some finding it extremely powerful indeed, while others, more in line with McKenna's suggested stance, finding it not strong enough to be worth getting worked up over. Given just how extremely powerful those on pro-AI side of this divide (by which I don't mean that those on the "pro" side in the end endorse AI, but just that they find it strong enough to bother over) report finding that premise, plus, I suppose, given how thoughtful and reasonable these philosophers are, I would have thought we already have a pretty good case based just on that that there's something here worth accounting for: even if not everyone is subject to it, at least on reflection, there is evidently an extremely powerful appearance of

ignorance popping up in a very surprising place (as the weak link in a two-premise argument to a bizarrely counterintuitive conclusion). This appearance evidently *can be* extremely powerful indeed, even if it isn't always so. One doesn't have to be overly curious to wonder why.

But one might worry that those who find AI so gripping were following some kind of intellectual fashion or group-think in finding it so. It was largely due to that worry that, quite a few years ago, I started asking introductory philosophy students, in the first meeting of an introductory level philosophy class, whether they thought they knew that they weren't BIVs, based just on a fairly minimal description of what BIVs are, and with no idea what answer I favored (which probably wouldn't have helped them with this yes-or-no question, anyway, as my view here is a bit complicated), in the way reported at (2017: 55-56). And as I report, the results were shockingly strong, in favor of the skeptic (in favor of the answer that one does not know that one isn't a BIV):

Well over 80 percent of students (and not just of those who voted) answered that they didn't know, and in each case, less than five (often in classes of about one hundred students, which is what this course drew for about three of those years, at least for its opening meeting during Yale's "shopping period") voted that they did know. (So there were a few abstainers.) (2017: 56)

I then repeated this experiment over the years, always ending up with similar surprisingly strong results. As I remark in the book: "I suspected the results were really about as strong as you can hope for from undergraduate classes on just about any question—including whether torturing and killing babies for the fun of it is wrong" (2017: 56). And remember, these are results obtained while I strove to be as neutral as possible on the matter. I'd be tempted to say something like "Imagine what the results might be if the matter were being pressed by an actual skeptic, trying to push for the negative verdict!"—but for the fact that it would be hard for results to get much

stronger than the ones that were actually obtained. Now, it is true, as I also report (at 2017: 56-57; and then in more depth in Appendix B, at 2017: 265-267), you get very different, rather skeptic-unfriendly, results if you ask the question in a different setting and in a different way. I venture some thoughts as to why there is this marked difference (at 2017: 57), and this seems a great topic for others to look into as well. (If you do so, and part of your looking involves asking folks what they think about the matter, it very much helps if you note exactly what you ask them (on this, see 2017: 57-58), and how you ask, so we can better compare notes.) But for now, you can put the matter this way (which comes close to how I put it at 2017: 59): That you *can* get results so strongly in favor the weak link premise of a two-premise argument to such a bizarrely counter-intuitive conclusion as that you don't even know that you have hands in the (quite boring and non-extreme) way it turns out you can is quite remarkable, and well worth accounting for, even if there are other ways of asking that return very different results. I mean, who would have thought you could? I certainly didn't. (I was genuinely shocked by the results the first time I tried asking introductory students, and even expected them to be very different the second time, only to be surprised that they were not. On subsequent tries, I suppose I was no longer surprised, but still found the results quite remarkable.)

McKenna's second line of critique (well, second in the order that I'm addressing them; first in the order he presents them) is his argument that "we don't need a contextualist semantics to account for the appearance of ignorance" (to the extent that there is an appearance here worth accounting for, I suppose). This because McKenna doesn't see why my contextualist solution, as presented in 1995/Chapter One, and further explained in the rest of the book, with some element of epistemic conservatism creeping in in Chapter Seven to start to address questions of

justification, can't be replaced by Sosa's safety-based solution, combined with epistemic conservatism.

So, first, a word about epistemic conservatism. Philosophically speaking, I was brought up on that stuff. I was an undergraduate at Calvin College when Plantinga and Wolterstorff's (ed.) book of epistemic conservatism, *Faith and Rationality* (Plantinga and Wolterstorff 1983) hung heavy in the air, and my professor for epistemology was Wolterstorff himself, whose own entry in that book (Wolterstorff 1983) was about as pure an expression of epistemic conservatism as you're likely to find. And my conservatism remained a strong interest as I started graduate school, to the extent that my first publication (DeRose 1989), on that late, great conservative, Thomas Reid, served for me as my way of exploring the conservative approach to epistemology and skepticism. But as I continued studying, I became obsessed with the basic argument by skeptical hypothesis (AI, as I later came to refer to that argument), as opposed to evidentialist arguments for skepticism (replying to which is the happiest home for conservatism), and also with the idea that a safety-based partial account of knowledge (though I wouldn't have been using such a label for it back at the time in question) could provide the explanation needed to solve the puzzle that AI confronts us with. After I finally published my solution, I think I came to be thought of, at least by some, as some kind of raving externalist about knowledge, where the truth was just that I thought an externalist feature of knowledge provided the solution to the AI puzzle. I try to explain this in the current book, and in Chapter Seven, my conservative ideas about epistemic justification start showing through a bit. I realize that as my picture of knowledge currently stands, it may now look like a disjointed cobbling together of two very disparate elements—a very externalist-looking double-safety condition together with an internalist, conservative account of justification that is also a condition on knowledge—but I really do think the two elements fit together well, or will do so,

once the conservative picture of justification is developed in the way I think it should be. One important commonality, important to the topic of skepticism, will be this: that even for someone with no fancy argument to the conclusion that they are not a BIV, and possessing no effective evidence for the claim that they are not, that they are not a BIV will be one of the things they know best (if I may speak that way), both in terms of how safe their belief in that is, and in terms of how justified they are in holding it.

But I hope I will be forgiven for putting off further development of the conservative approach to epistemic justification to future work.¹ As McKenna realizes, conservatism alone isn't what's needed here; the comparison before us turns on what conservatism is combined with, and McKenna is wondering: why not replace my contextualist account with the safety-based account Sosa offers, primarily in (Sosa 1999), on which, as McKenna nicely summarizes it:

The basic thought is that, while knowledge requires safety rather than sensitivity, it is very easy to confuse safety with sensitivity. It is therefore not difficult to see how a skilled skeptic could lead us to deny that we know the denial of skeptical hypotheses (or for that matter to deny that we know anything at all) by appealing to reasons for thinking that our beliefs are insensitive.

As McKenna realizes, I did address the comparison between Sosa's solution and my own at considerable length in (DeRose 2004), but he is wondering why in a book-length treatment, I don't at least acknowledge the issue here, and give some indication of how I think it could be tackled.

¹ My current big project is a book on the problem of evil. My plans (though you never know how things turn out, so, you know, *Deo volente*, and all that) are for my next big project (I often mix in some littler things to work on concurrently with my big projects) after that to be on epistemic conservatism. For some idea of how some of my thinking on this goes, beyond the glimpses afforded in Chapter Seven of this book, and how my differs from what I imbibed all those years ago at Calvin College, see (DeRose 2005), which does come up, if only to be set aside, in the current book (2017: 225-226).

In the end, the omission *may* have been a mistake on my part, but I can explain why I went that way.

I have always been in some good sense, an “insensitivity theorist” in that I’ve thought that the insensitivity of some key beliefs plays a starring role in explaining why those beliefs don’t seem to constitute knowledge. And Sosa has been a critic of me as such, utilizing counterexamples to the generalization that we tend to think to think that insensitive beliefs don’t amount to knowledge. This always led me to wonder why Sosa thought it was that we don’t take ourselves to know such things as that we’re not brains in vats. As we’ll see later, one very important insight Sosa had here was that the intuitive situation here is not so simple, and the complexity should be accounted for. Still, there was at least a discernible intuitive push towards some judgments of no-knowledge that I accounted for by appeal to the insensitivity of the beliefs involved, and it wasn’t clear to me where Sosa thought that push was coming from. And then Sosa advanced the explanation well-summarized by McKenna in the quotation indented above, and, well, yes, that did seem to have potential for explaining the matter.

But wait! Didn’t that explanation also make Sosa an insensitivity theorist of a sort—and indeed, of very much the same sort I am? Because I am not a *direct* insensitivity theorist, as that term is explained at (DeRose 2004: 24) and also in the current book at (2017: 179). Like Sosa (as he’s summarized by McKenna above), I appeal to the relevant belief’s insensitivity to explain why we think they don’t amount to knowledge, but also like Sosa, I don’t think this is so because knowledge has a sensitivity condition. This is why we’re not direct sensitivity theorists. So Sosa, and I, and direct insensitivity theorists all seem to subscribe to something like this “Insensitivity Account” (IA): in the relevant cases, we all explain why S seems not to know that p by means of claims like:

1. S's belief that p is insensitive, and
2. We have some at least fairly general—though perhaps not exceptionless—tendency to judge that insensitive beliefs are not knowledge.

Direct sensitivity theorists further claim that (2) holds because:

- a. Sensitivity is a necessary condition for knowledge

But Sosa and I are indirect insensitivity theorists, because we think (2) holds (or at least I don't see how to make Sosa's explanation fly without using *something like* (2)) even though knowledge does not have a sensitivity condition. We each give some other basis for (2). And in Sosa's case, that basis is well-expressed in McKenna's summary of his account: knowledge has a safety condition, but we easily confuse safety and sensitivity. One important upshot of all this, I think, is that, on the face of things, at least, Sosa's own account is hurt by his counterexamples to (2) as much as my account is! Well, actually, (2), as I've formulated it, doesn't really have "counterexamples" to it, as I've worded it so it not only allows for, but positively anticipates, exceptions to the generalization that it is advancing. Sosa's "counterexamples" do hurt my account, to the extent that they do, by showing that the generalization I appeal to has exceptions, while it would be better to explain by means of a generalization that doesn't have as many—or perhaps that doesn't have any—exceptions. But since Sosa's account seems to make use of something like (2) as well, it's hard to see why it isn't as damaged by the "counterexamples" as my account is. Here, as elsewhere, it seems as incumbent on Sosa to get into the weeds and explain various things in some detail as it is on me. Note the comparative nature of that statement. Like me, Sosa can improve his account of the no-knowledge judgments we've been considering by means of a more exact statement of the generalization he uses. (And maybe he'd like to use some of the moves I go

through in Chapter Six to handle counterexamples, toward that end.) But also like me, he could also argue that his explanation is a good one, even though/even if there are exceptions to the generalization he uses. What he can't legitimately do, at least without providing some basis for doing it, is fire off counterexamples and then take them to be my problem, but not his. At least, that's the argument I pushed in (2004). But I now wonder if I was over-reading Sosa's commitment to accounting for any strong drive toward judging that insensitive are not knowledge?

But before explaining that question, let me give an example of the kind of thing that I think is better explained on my account than on Sosa's. My account of why we tend to think that insensitive beliefs don't amount to knowledge (my account of (2), above) is based on my contextualist double-safety account of knowledge, and my "Rule of Sensitivity." That Rule is for raising the standards for what counts as "knowledge," and, given that we sometimes have conversational purposes that are advanced by using unusually high standards for knowledge, "the usefulness of some device like the Rule of Sensitivity is apparent," as I put it at (2017: 210), and it makes sense that we would have such a device, by which we could indicate with some precision just how high we intend to raise the standards for "knowledge," in the way explained at (2017: 210-211). But like other rules for "changing the conversational score," there should be situations in which, for a variety of reasons, the rule doesn't operate. And I've long pressed the examples that appear in section 6.8 of the current book (2017: 188-192)—these go back to my dissertation (1990: 219-225)—as cases where one can and naturally would truthfully claim that someone knows that some skeptical hypothesis is false, even though their belief that it is false is "really" insensitive (and not just insensitive by some preliminary understanding of sensitivity). These claims are very different from the bragging engaged in by my "boastful zoologist" at (2017: 9). The best sense we can make of that braggart's claim is that he is utilizing the Rule of Sensitivity

to raise the standards for knowledge by his positive assertion, and to thereby claim to meet, not just ordinary, but somewhat unusually high standards for knowledge, with respect to his knowing that the animals in the zebra cage really are zebras. That's just the kind of thing the Rule of Sensitivity is there to facilitate. But my speakers in section 6.8 clearly have a different conversational purpose in play: not to ascribe unusually elevated knowledge of not-H to the subject they're talking about, but to (favorably) *compare* that subject's standardly good epistemic position with respect to not-H² to the unusually and deplorably bad epistemic position that some poor unfortunate other character in the story is in with respect to not-H, and to thereby report that the subject is in the standardly good epistemic position with respect to not-H (as opposed to the poor unfortunate, who is not). So, here is the example involving "knowledge" that some zebras are not cleverly painted mules:

Suppose a group of friends is meeting by the zebra cage at the local zoo, and as two of the friends, Nico and Lou, who are a bit late, approach from a distance, one who is already at the meeting place, Andy, devises a practical joke to play on them: The rest of the friends will try to get Nico and Lou to falsely believe that the zebras are cleverly painted mules. When Nico and Lou finally arrive, the friends around them start saying things like, "What a great paint job! I would have never guessed that they're just painted mules," and, "When I heard of this painted mules scam, I never thought the zoo would be able to get away with it so easily." After Andy walks away to buy several stuffed zebras from a nearby concession stand, Nico catches someone snickering, and

² And it is a feature of my position that we are in standardly good epistemic positions with respect to the falsehood of typical skeptical hypotheses: our usual tendency to think we don't know that they're false is not due to our being in such a weak position with respect to that falsehood, but is rather due to the usual operation of the Rule of Sensitivity, as explained at, e.g., (2017: 29)—though at that point I'm giving the explanation as if the Rule is always operative and always carries the day.

has to be let in on the joke, but Lou is completely taken in by it. So, when Andy returns, one of the friends pulls him aside to apprise him of the situation, explaining: “Nico knows that they’re not painted mules. Sorry, we had to tell her. But Lou totally bought it!” Here Nico, like all her friends, has not performed any of the investigations that would make her believe that the animals aren’t cleverly painted mules sensitive, and as a result, her belief is (“really”) insensitive. Yet here it seems, at least to me, that “knowledge” that the animals are not painted mules is truthfully attributed to her. Given the conversational circumstances, the assertion that “Nico knows that they’re not painted mules” does not have the effect of raising the epistemic standards to the unusually high level at which Nico’s belief in that proposition must be sensitive to count as “knowledge.” The point our speaker fairly clearly means to be making is not that Nico has the kind of *fairly* exalted knowledge of the fact that the animals are not painted mules that some special investigation (perhaps, as has been suggested (Stine 1976: 252), one involving paint remover) or special knowledge (as in the case of my “boastful zoologist” in SSP, Section. 1.3) would allow her to have, but just that she has the more humble (and strangely difficult to claim, in *most* conversational circumstances) “knowledge” of such facts that most trusting zoo-goers at reputable zoos have—but that Lou here lacks. (2017: 191)

Here is an example which is an exception to the operation of my Rule of Sensitivity, but I found such exceptions specifically by construing the Rule as a mechanism for changing the conversational score, and then trying to come up with conversational situations where I would expect the Rule not to operate. As such, it all seems to make good sense: here’s the Rule; here’s the purpose it serves; here’s why ascribing knowledge of the falsehood of a skeptical hypothesis tends to engage that Rule in order to serve that purpose; but here’s a conversational situation where

it would be fairly clear that “knowledge” of not-H is not being ascribed in a way that engages that purpose, but for another reason, and so we can expect the Rule not to operate. (As I would expect, a similar case produces a situation in which it is natural to, and seems true to, ascribe to someone “knowledge” that they have lost a lottery, even though their belief that they lost is really insensitive; see (2017: 157-159)). By contrast, I don’t know how to explain this on Sosa’s confusion account. Here our speaker says that Nico knows that the animals are not cleverly painted mules, and so at least in some sense would be *thinking* that Nico know this, though Nico’s belief is insensitive. But if we tend to think that insensitive beliefs are not knowledge for the reason Sosa gives, why doesn’t our speaker here confuse safety with sensitivity, and think that Nico doesn’t know? Is there something about the situation that works against that confusion setting in here? I don’t see how to make this work.

And in fairness, I suspect that Sosa doesn’t try hard to make such things work because he doesn’t really believe in the phenomenon being explained here, at least in the way I do. (Here we connect up with the first main criticism of McKenna’s that I responded to.) On my view, the intuitions that insensitive beliefs are not knowledge can be very strong indeed (even if a bit “flighty,” as I put it) in lottery cases, and also in (and in a less “flighty” way) cases of the non-obtaining of quite mild skeptical hypotheses. And when it comes to beliefs that radical, remote skeptical hypotheses do not obtain, even here the intuitions that these beliefs don’t amount to knowledge *can* be quite strong indeed. For me, all this not only calls for, but cries out for explanation. And where there are cases where the expected no-knowledge intuition doesn’t materialize, I really want to know, if I can, why that exception occurs. But I think that for Sosa, the intuitions just aren’t that strong, and are thought of as a bit more of a mess, and so he is happy with explaining the drive (or maybe better, for him: the mild “push”), insofar as it is there, toward

judging insensitive beliefs to be cases of no-knowledge in the way we've looked at, and in a way that doesn't get concerned with accounting for exceptions. (It's like I imagine him thinking: "There's an exception? Well, you know, it wasn't that strong of a tendency, anyway.")

I'd like to add that while I think it has been, and remains, the case that Sosa and I have quite different perspectives on just what needs explaining in dealing with skepticism, it has been extremely valuable for me to encounter, wrestle with, and learn from his view. In particular, it was immensely valuable for me to encounter this passage, also from (Sosa 1999):

Consider, moreover, the need to explain how the skeptic's premise—that one does not know oneself not to be radically misled, etc.—is as plausible as it is. That requirement must be balanced by an equally relevant and stringent requirement: namely, that one explain how that premise is as implausible as it is. To many of us it just does not seem so uniformly plausible that one cannot be said correctly to know that one is not at this very moment being fed experiences while envatted. So the explanatory requirement is in fact rather more complex than might seem at first.

(1999: 147)

My account of the intuitive complexity involved here, explained in section 2.10 of the current book, is I think one of the most important aspects of my treatment of skepticism, and would not exist if I had not benefitted from Sosa's work.

But because, as I see it, Sosa wasn't involved in trying to give the kind of account I thought was needed, it didn't seem right to line up the various things I thought my account could explain better than his. Rather, I tried to make the case as best I could that there really was an important drive toward judgments of no-knowledge that needed to be accounted for in a serious way, and

then give that account (including explaining the intuitive complexity that Sosa pointed out), leaving (DeRose 2004) to stand as my case for my account over Sosa's, insofar as one is inclined to take Sosa as a serious, if indirect (to use my terminology), insensitivity theorist.

2. Reply to Blome-Tillman: On Counter-Examples

I'm not sure what to make of the issues raised by Blome-Tillman in the first paragraph of his section 3, on "Some Potential Problems." These are problems, or at least apparent problems, that I deal with in this two-book series. In his next paragraph, Blome-Tillman writes, "DeRose himself is, of course, aware of these difficulties," but, aside from relaying a feature of my position relevant to one of the worries, doesn't point to my answers. So I think I will here direct readers to the relevant parts of the books, adding a few words on what they will find there. On the supposed "elusiveness of knowledge": I believe contextualism does not make knowledge elusive in the relevant way, and I thought how Lewis presented the issue was a mistake. Lewis disagreed, in a spirited way, and readers may enjoy a passage from a letter Lewis sent me on the matter in which he scolds me (as I think it's fair to say) for, well, scolding him (also fair, I suppose) for how he does this, which I quote in Volume 1 at (2009: 217). The issue of elusiveness—both the strategic difference (as I think it's fair to call it) I had with Lewis, as well as the substantive issue of whether there's a real problem here, which I fight out a bit with Dretske—is worked over in the last four sections of Chapter 6 of Volume 1, at (2009: 212-225).

I don't feel *certain* about what Blome-Tillman thinks is the status of the charge that "DeRose's account entails the odd result that assertions of the sentences 'I know that I'm not merely dreaming/that I am not a brain-in-a-vat' are self-undermining: they cannot ever be made

truthfully,” but I deny that my account has this result, and make the case for this in the current Volume in section 4.16, at (2017: 121-123), even closing that section by myself making a claim that “I [do] know that I’m not a BIV” that I claim is true. (Readers may enjoy the story, at the start of this section, of Rogers Albritton, who was my dissertation advisor, raising this problem while I was a graduate student.)

What I think is going on is that Blome-Tillman holds that my moves here, which involve my exploiting some flexibility in how the epistemic standards are set, don’t fit with my Rule of Sensitivity; as he writes, “[I]t is noteworthy that it is quite unclear how such a significantly more flexible (and attractive) view about the variability of epistemic standards can be reconciled with the rather rigid Rule of Sensitivity.” The answer is that they go together well because the Rule of Sensitivity is not, and never was, so rigid, in terms of when it does and when it does not win out over other conversational forces, resulting in the epistemic standards actually rising in the way it points to. When I introduce the Rule in Section 1.12 of “Solving the Skeptical Problem,” this takes place against the background, set in section 1.2 (see the “second important point,” at 2017: 4-5), that I’m assuming, just for ease of exposition, a “skeptic-friendly” version of contextualism on which the Rule always carries the day in raising the standards. It is thus understandable how it could appear to be rigid in the relevant respect. But this really was an assumption only made for ease of exposition, and, in fact, as a Rule for changing the “conversation score,” it will often come into conflict with other conversational forces, and will not always win. In the current book, I explain repeatedly how the Rule thus really only posits a *tendency* for the raising of the standards (including at 2017: 122, in the section where I address the supposed inexpressibility of my position), and defend how it can achieve its explanatory purposes nonetheless.

On the Rule of Sensitivity and my Bank Cases: I believe that when I first formulated those cases, I considered having my wife say “You don’t know that the bank hasn’t changed its hours in the last two weeks,” which would have rendered her closer to in effect using the “Argument from Ignorance” in that case, and would have made her claim fit the triggering condition of the Rule of Sensitivity, even on the preliminary formulation of that Rule that I use in “Solving the Skeptical Problem” (and in Chapter 3 of my dissertation (DeRose 1990), the precursor to that paper), but decided not to do that, specifically in order to gain a *bit* of “distance” between how she raises the standards for knowledge and how skeptics at least threaten to raise the standards in presenting AI, and specifically to illustrate that we’re dealing with a broader phenomenon than our treatment of AI-inspired skepticism might indicate. (How much “distance” I so gained is quite questionable, though, as I still have her raising what is in effect a (mild, local) “skeptical hypothesis,” even though she doesn’t explicitly charge, but only insinuates, that I don’t know that it doesn’t obtain.) Thus, when I introduce the Rule to help solve the AI puzzle, I start by saying, “Although a more general formulation of this rule is desirable, I will here state it in such a way that it applies only to attributions (and denials) of knowledge, since such applications are what’s needed to address the present puzzle” (2017: 27), acknowledging that something like the Rule is likely of more general application—and of more general importance, I suppose. I guess I don’t see this as a problem.

Blome-Tillman’s main criticism involves a counterexample that he raised in his (2009), and presses again here; it is his interesting

- (1) Smith knows both that no material object can travel faster than the speed of light and that the animals in the pen are zebras.

He explains the trouble it causes for me:

Intuitively, there are possible contexts in which (1) can be asserted truly, and any contextualist theory should license this result. On DeRose's account, however, (1) can never be truly asserted. That is so because the first conjunct of the 'knowledge'-ascription expands the sphere of worlds throughout which our belief must be safe to an extraordinarily far away point in modal space—namely, all the way out to the closest world in which material objects can travel faster than the speed of light, and thus to a world in which the laws of nature differ dramatically from our actual laws of nature. Within the safety sphere that stretches that far out, however, are worlds in which the animals in the pen are cleverly disguised mules. Worlds in which the animals are cleverly disguised mules are, after all, significantly more similar to actuality than worlds in which material objects can travel faster than the speed of light. Thus, the first conjunct of (1) raises the epistemic standards to a level that makes it impossible to satisfy the entire, conjunctive knowledge ascription. Knowledge of (1) cannot be truthfully asserted.

I certainly agree with the intuition Blome-Tillman reports at the start of the above quotation. What I'm worried about is whether, in the relevant sense, we have to travel to such an “extraordinarily far away point in modal space” to get to the nearest world in which Smith's first belief is false—but that “in the relevant sense” needs explaining. In the end, it's not so much that I can see that the nearest worlds in which some material objects can travel faster than the speed of light is really quite close in the relevant sense as that I don't share, and don't see the basis for sharing, Blome-Tillman's confidence that they're so extraordinarily distant.³

³ What's directly at issue is whether the closest such worlds are close enough that they don't allow for still closer worlds in which the animals in the pen are not zebras, but one confidently believes they are. Now, how close that is may be up for grabs a bit. As epistemologists (or at least those who move in the circles that I do) are by now all aware, there have been zoos that actually have used painted mules (though how cleverly they're painted is

I have a lot to say in the book about philosophical methodology, especially as that concerns counterexamples: see especially sections 6.1, 6.3, 6.5, and, especially relevant here, sections 7.3-7.4. On the methodology I follow, for me, a lot of proposed counterexamples to my account instead turn into opportunities to more precisely understand the notions of safety and sensitivity, as these figure in epistemic evaluations, and especially how we should understand the notion of the “closeness” of possible worlds, in the epistemically relevant sense. While I suspect Blome-Tillman might not like that methodology, he doesn’t really engage with it or my defense of it, so what I will do here explain the relevant notion, then turn to his example.

Let’s start with Nozick’s grandmother case, and then my now quite old treatment of it:

A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn’t know he is well (or at least ambulatory) when she sees him. (Nozick 1981: 179)

This is a nice case by which Nozick challenges a preliminary version of his theory of knowledge. To let it do its work, we have to (follow clues Nozick gives above and) understand “well” so that one counts instead as unwell only if one is very unwell indeed. This is to nail down the clear intuition that the grandmother does know that her grandson is “well.” But with that understanding in place, the intuition that the grandmother, in the circumstances we naturally imagine her as being in, does know that her grandson is “well” is quite strong indeed. However, she seems to fail

questionable: when pictures are included, I’m usually left doubtful that even I, zebra novice that I am, would be really fooled by these creatures), as we are occasionally alerted to amusing blurbs in magazines and the like that report such shenanigans actually taking place. But I take it it is the case that (and if not, I should here pretend that) this deplorable practice isn’t widespread and threatening enough to block our knowledge-by-ordinary-standards, while in reputable zoos, that we’re seeing really seeing zebras in the zebra cage, and thus that the nearest worlds in which the animals before one in the standard setting in a reputable zoo are deceptive enough fakes that they would have completely fooled you are at least fairly distant.

Nozick's third condition for knowledge, as it stood at that point in his exposition, which would require that

If the grandson had not been well, the grandmother would not have believed he was well

Nozick responded by modifying his third condition, inserting a reference to the subject's method of belief formation into that condition's formulation. But back in 1995,⁴ I wondered whether we might instead understand the simple (in formulation) conditional above in a particular way, on which it *is* satisfied in this case:

On one way of evaluating that simple conditional, it seems that if the grandson were not well, the grandmother would *not* believe he was well. After all, she's looking right at him! The standard possible-worlds semantics for counterfactual conditionals can illuminate what's going on here. When one searches for the possible worlds most similar to the actual world in which the grandson is not well, the respects in which the possible worlds are to resemble the actual world is a highly context-sensitive matter. Especially where the context focuses one's attention on the grandmother and her cognitive and recognitional abilities, one *can* place heavy weight upon similarity with respect to the method she is using to arrive at her belief, and then it can seem that in the closest world in which the grandson is not well, she's looking right at him and seeing that he's not well, and so does *not* believe he is well. On this way of evaluating the conditional, the grandmother *does* satisfy even the initial formulation of Nozick's third condition, and she's no counterexample to the generalization utilized by SCA.

⁴ The passage I'm about to quote is at (2017: 16) of the current book, but is unchanged from (1995: 21).

As the end of that quotation shows, I was, and am, concerned with the “SCA,” or “Subjunctive Conditionals Account,” on which we explain certain appearances of ignorance by appeal to the generalization that we tend to think subjects don’t know things when their beliefs in those things are insensitive, a generalization that Blome-Tillman’s case also calls into question.

Let’s move to another case, to serve as background, this one based on an example of Jonathan Vogel’s (1987: 206), but formulating it as follows:

The case of the non-melted ice: Quite a few hours ago, you left a glass with a little bit of ice in it out in your screened-in porch. As you realize, it’s been a very hot few hours out on your porch, so you feel quite certain that the ice all completely melted a long time ago. However, the glass is opaque, so, though you’ve been able to see the glass sitting there undisturbed this whole time, you weren’t able to see the melted ice. But we (not you in the example, but us talking about the example) can stipulate that you (in the example) are right: the ice of course has completely melted, and what’s more, nothing funny has been going on: everything is quite boringly as you’d naturally imagine it from my boring description of events.

I take it an intuition about this case is this (“you” here is you-in-the-example):

If the ice in the glass hadn’t melted, you (still) would have believed it had melted

At least, that’s how I admit I evaluate it when I just consider it in the way I’m offhand most naturally inclined to understand the conditional—and I gather that’s how most respond. But it’s worth pausing to ask: How am I understanding that conditional in this offhand way, when I judge it to be true? To put this in terms of possible worlds, how am I imagining the nearest world(s) in which the ice hasn’t melted? And here I must report that I’m not imagining that some neighbor

has slipped in unnoticed and put into the glass some advanced (I suppose battery powered) cooling device, which has kept the ice from melting, or anything like that. I'm just imagining that the ice has been sitting out there in the heat all this time, but has just, I suppose by some local (but temporally quite sustained) miracle, failed to melt, and aside from that weird failure to melt, everything has been happening normally, leaving you to judge very reasonably, but as it turns out mistakenly, that the ice has melted. That bit about everything else happening normally is perhaps worth stressing. I don't think, "Wow, if something as wild as a miraculous failure to melt were happening on the porch, who knows what other crazy things would be happening? All bets are off! I have no idea what would have been going on if that wild thing had happened!" No, when I find the conditional indented above straightforwardly true, I seem to just change that one fact about the situation, not worry about my introducing a rather sustained miracle into the world to do so, and suppose that the world then and otherwise proceeds to operate as normal. Now, all this doesn't match up with the offhand judgment of similarity I would make if I were just asked about which world(s) in which the ice hasn't melted are the ones that most resemble the actual world. And it seems a good bet that it wouldn't line up with Blome-Tillman's similarity judgment, either, what with the sustained miracle occurring in what I'm taking to be the closest no-melt worlds, while I suppose it would be easy enough to secure a world in which the ice hasn't melted that contains no miracles. But it is in keeping with David Lewis's account of the standard reading of counterfactual conditionals, as he thought the nearest world in the relevant sense in which the antecedents of those conditionals are true are standardly miracles away from the actual world. As he memorably remarks in one place where he's comparing a couple of possible worlds in terms of their similarity to the actual world, as that similarity is measured in the way relevant to our standard understanding of counterfactuals: "The lesson we learn by comparing w_1 , and w_2 is that under the similarity

relation we seek, a lot of perfect match of particular fact is worth a little miracle” (1979: 469). And he issues this memorable “word of warning” (that I cite in the book at 2017: 208):

It is all too easy to make offhand similarity judgments and then assume that they will do for all purposes. But if we respect the extreme shiftiness and context-dependence of similarity, we will not set much store by offhand judgments. We will be prepared to distinguish between the similarity relations that guide our offhand explicit judgments and those that govern our counterfactuals in various contexts. (1979: 466)

In keeping with Lewis’s advice, I have long thought (and this is showing already in my old reaction to Nozick’s grandmother case) that there is a particular way of construing the relative closeness of worlds that is relevant to epistemic evaluations and that often diverges, not only from the offhand judgments of comparative similarity we would make, as already happens with the standard understanding of conditionals, but also sometimes with how we standardly understand counterfactuals for other purposes. In his (2009), Blome-Tillman considers something like this move, but he puts it in terms of taking me to be “not talking about overall closeness to [the actual world] but rather about closeness to [the actual world] in some particular respect” (2009: 390). However, following Lewis’s treatment, I think of the epistemically relevant understanding of conditionals to be governed by an overall similarity relation—just one that differs from the overall similarity relation we use in the standard understanding of counterfactuals. What’s important is that it be a good way, even if not the standard way, of understanding the relevant conditionals in ordinary English, and that it provide a reading of the “insensitivity conditional” (“You would have believed that *p* even if *p* had been false”) that is true of you in the standard no-knowledge cases as they’re standardly construed (where you intuitively fail to know that you’ve lost the lottery or that you’re not seeing cleverly painted mules), and that well explains why you seem not to know in

those cases, and that ceases to be true of you where those cases are altered so as to make it seem that you do now know (so, like, after you have heard the winning numbers of the lottery announced, and see that they don't match the numbers on your ticket, or after you've tried paint remover on the poor beasts in the zebra case at the zoo). I would love to present a general and fairly detailed account of this epistemically relevant way of construing similarity, of the form of Lewis's account of the similarity relation governing the standard meaning of the counterfactuals at (1979: 472)—you know, in “It is of the first importance to....It is of the second importance to....[etc.] It is of little to no importance to....” terms. But, sadly, I'm not ready to do that yet. But I can tell you what I think in the relevant sense are the closest worlds in which the ice hasn't melted in the Vogel-ish case, before then venturing some tentative thoughts about what in that relevant sense are the closest worlds in which some material objects can travel faster than the speed of light, for Blome-Tillman's example.

And I can say a bit about what intuitively guides me to the worlds in question as the closest no-melt worlds on this epistemically relevant way of measuring, and I can also first say something about an important way that the epistemically relevant way of understanding of counterfactuals differs from standard way of taking them. While we can understand counterfactuals in a backtracking way—where we suppose the antecedent of the conditional, and then reason backward in the temporal/causal order to what would have been the case so that A would have been true, in order to find the closest A-world, and then forward on to C: “If A had been the case, that would have been because...., and then C would have been the case”—we standardly don't. But going back to some of Nozick's interesting musings in *Philosophical Explanations* (see 1981: 223n), it has long been suspected that the epistemically relevant way of understanding counterfactuals *would* allow backtracking.

Now, to the unmelted ice. As I indicated back in my old treatment of Nozick's grandmother case, I follow Nozick in stressing the importance of seeking worlds in which the subject uses the same method of belief formation as she does in the actual world. In implementing that, I think we should pay special attention to what the subject is "going by" in reaching her belief. In our ice cubes case, she is going by an approximate idea of how hot it is out on the porch and also by an approximate idea of how long the glass that started out with ice in it has been out there. But we don't construe our subject's method so that it doesn't count as using the same method if it yields different results. In fact, we should be very open to worlds where using the same method would issue very different results. With all that, and an openness to backtracking in mind, I think the closest worlds in which the ice hasn't melted, in the epistemically relevant sense, are worlds in which it hasn't melted because it's very cold outside and/or the glass has only been out there for a very short time. And in such worlds, you don't believe the ice has melted. So, in the epistemically relevant sense

If the ice in the glass hadn't melted, you would have *not* have believed that it had melted

Your belief is not insensitive.⁵ It's worth noting here that while our epistemically relevant way of construing similarity is different from how we construe similarity in our standard understanding of conditionals, this difference may here be moving us closer to how we construe similarity in our offhand explicit similarity judgments, given that I'm not supposing there are any sustained miracles going on in my epistemically closest no-melt worlds. Of course, one might wonder just how we get from our hot actual (actual in the story) world to, say, the freezing cold no-melt world.

⁵ Well, your belief that the ice has melted isn't insensitive. But this belief, if you have it, *is* insensitive: that it's not the case that by some sustained but local little miracle, the ice has failed to melt.

But the short answer is that, as I evaluate the conditional in this non-standard, but still natural, I think, epistemic way, I just don't worry about that. If someone explicitly asked me what you would have believed if you had left the glass out there hours ago, but if it had been freezing cold outside, I don't think, "Oh, my, what a wild, remote possibility we're considering here! Who knows?" I just think, "Well, if it had been so damn cold out, you probably wouldn't have believed that the ice had melted—or at least wouldn't have been so damn sure about it,"⁶ without any sense that I was considering some very wild, remote possibility. Similarly, when I understand the conditional indented above in the epistemically relevant way, I judge that conditional true without any sense that I'm contemplating what's happening in some remote, weird region of modal space.

With all that in mind, as we turn back to Blome-Tillman's example, it's far from clear to me that in the relevant sense—the sense I hope you have been getting a sense of—the closest worlds in which some material objects can travel faster than the speed of light are so fantastically distant. In some good sense, such worlds may be very distant indeed. When thinking in those terms, one might be moved to exclaim: "Good heavens, to get from here to there you have to *CHANGE THE LAWS OF NATURE!!*" But remember Lewis's word of warning! Consider someone who takes themselves to know that no material object can travel faster than the speed of light, based on, say, their science textbook in school saying so, and seeming to be taking it as a fairly settled matter, and folks around them whom they trust on such matters seeming likewise confident. If such a person, while thinking in terms appropriate to epistemic evaluations, considers what would have been the case if that fact had not been so, and they find themselves trusting that if it had not been so, then the science smarties around them wouldn't be telling them so, at least with such authority

⁶ On the potential importance of this last bit, see the sections of the book on "Degree of Belief Sensitivity," sections 6.11-6.13, at (2017: 196-201).

and confidence, and so they themselves wouldn't be so sure that it's so, I don't think they would get any sense that they're contemplating what, in the relevant way, is some wild, remote scenario. So I'm not so sure about the "distance" involved here.⁷

All that said, modal epistemic notions like safety and sensitivity seem most at home in evaluating beliefs in things that, in metaphysical terms, easily could have been different, as opposed to beliefs in things like mathematical truths or laws of nature, and so I feel quite tentative about matters here. (This is connected with how, at key points in the discussion above, I find myself relying on the likes of a sense that I'm not considering "some wild, remote scenario," instead of finding guidance in judgments about the truth-values of counterfactuals that I feel are quite solid.) As I indicate in the book at (2017: 213), in connection with some such cases, I seriously contemplate taking the notion of safety in the direction David Manley has proposed (Manley 2007), on which, roughly, safety is measured by the nearest worlds in which one goes wrong *about some counterpart thought*, but, to go beyond what is reported in the book, I am having trouble figuring out how to construe sensitivity so that it will correlate with that revised notion of safety. (But I'm also wondering whether I need to get sensitivity to so correlate in order for it to do the explanatory work I put it to, if knowledge is construed in revised safety terms.)

So I admit this is all rather unsettled for me. But for reasons for sticking with insensitivity accounts of important epistemic matters, see especially my arguments in section 6.5 (2017: 182-183).

⁷ Perhaps it will help if I add: while I'm not seeing that the worlds in which superluminal travel is possible are so distant in the relevant sense, things are very different when it comes to worlds where superluminal travel is possible but some demon messes with all our attempts to discover this and the demon makes it seem to us that such travel is impossible. Now *that's* a fantastically remote possibility in the relevant way!

3. Reply to Brendel: Disagreement, Retraction, and the Skeptic

I felt the strong urge to defend my views from the objections Elke Brendel lays out and explains in the fourth and fifth paragraphs of her Introductory section, for, though she then goes on to relate in her next paragraph that my book “addresses some of the above mentioned worries,” the answers are not explained, nor pointed to, leaving the objections hanging. But I resist doing so, to get to the problems Brendel focuses upon: contextualism’s supposed difficulties with disagreement and retraction. In a way, these objections are more germane to my first volume in this two-book series (DeRose 2009), and part of what will be involved in answering them is recalling things done in that older book, but Brendel does specifically relate these issues to disputes over skepticism, and I, too, will keep things relevant to the new book by being especially concerned here with disputes that have a skeptical bent.

There is a gap in my treatment of disagreement in Vol. 1. Following John MacFarlane, Brendel is worried about contextualism’s problem of “lost disagreement”: Does contextualism fail to secure genuine disagreement in cases where one speaker says someone “knows” something, another speaker says that same person doesn’t “know” that same thing, and it seems like the speakers are disagreeing with one another? My “Single Scoreboard Semantics” approach is supposed to secure genuine disagreement between such speakers, in ruling that the claims such speakers make are incompatible with one another (they can’t both be true), in some of the cases where it is most urgent to secure this result, because it is intuitively clear that the speakers are disagreeing with one another: cases where the speakers, perhaps in addition to indicating what standards for knowledge they intend to be using, also each indicate that they mean to be disagreeing with each other.

I worry that the scope of cases this approach can handle gets underestimated by some critics. One of the main upgrades in “Single Scoreboard Semantics” in its transition from an earlier journal article to a later chapter in Volume 1 is that I added, in section 12 (2009: 148-150), a treatment of “one-way disputes,” where, for instance, a speaker, S1, has said, “S knows,” and a later speaker, S2, in a different conversation in which, on contextual analysis, higher standards seem to be operative, disputes S1’s earlier claim, saying, “S1 was so wrong! S doesn’t know.” The sense that S2’s claim that “S doesn’t know” is incompatible with S1’s earlier that “S knows” can be pretty strong, and the “asymmetrical gap” analysis I suggest for such cases does secure the result that these claims can’t both be true, and it does so without holding the content of the earlier claim hostage to what is said later and in a different conversation. However, MacFarlane seems to miss this, as he writes: “However, this strategy is limited. It applies only to cases where the two disagreeing parties take themselves to be participating in a single conversation” (2014: 181, n. 5). And Brendel repeatedly writes that my approach is limited to “face-to-face communications.” But I don’t intend for it to be so limited. I set the discussion up with an imagined confrontation between Moore and skeptic discussed by Mark Richard (2004), which seems to be a face-to-face encounter. But what is supposed to trigger the single scoreboard approach is not that you have two disputants who are somehow face-to-face with one another, but (I guess what often goes along with that) two speakers who are indicating that they mean to be saying things in conflict with one another, and what triggers my approach to one-way disputes is a later speaker who indicates that they mean what they say to be in conflict with what was said in another conversation. Brendel worries about my approach making the truth conditions of knowledge ascriptions depend “not only on the speaker’s standards” but to “also depend on the knowledge standards of potential other speakers.” But it depends on those other speaker’s “personally indicated content,” as I put it, because the

speaker in question made it so depend, by indicating not only that they intend to be using such-and-such standards, but also that they intend to be contradicting the other speakers.

However, the approach is of limited scope—even if not as limited as it is sometimes taken to be. Brendel and MacFarlane have in mind, as trouble for contextualism, “disagreements” between speakers who are engaged in separate conversations, usually unaware of each other. This brings us into the issue I address in section 2 of Chapter 5 of Volume 1 (2009: 155-159): “The Objection [to Contextualism] from Judgements of Comparative Content.” But here’s where the big gap in treatment comes in. Contextualists certainly do think that it often happens that one speaker says that some subject “knows” something to be the case, and another speaker says that same subject “doesn’t know” that same fact to be the case, and both speakers are speaking truthfully, and so compatibly. In fact, my main argument for contextualism is driven by just such cases. But when I answer the “Objection from Judgments of Comparative Content,” arguing that when we look at the right pairs of test cases, the intuition that the two claims are incompatible with one another is at best extremely weak, my “right” case pairs are limited to the ones that have all the features I think make for the very best case pairs for arguing for contextualism. But that leaves my response to issues of “disagreement,” where we use that broadly to cover all cases where one speaker says someone “knows” and another that the subject “doesn’t know,” whether or not these speakers are in contact with one another or even know of each other’s existence, to the very extremes on either end: the cases in which it intuitively seems clearest that the claims by the speakers are incompatible with one another, on the one hand (where I do secure the incompatibility of the claims), and the cases where we are most inclined to take each claim to be true (where I argue there is no great intuitive cost to ruling that the claims are compatible), on the other.

But what of the vast middle ground? I suspect that the best case against contextualism from judgments of comparative content will be somewhat analogous to what I suggest as the best objection to contextualism from belief reports in Volume 1, Chapter 5, section 6 (2009: 166-170): that, while we (speakers of English) may not *always* be quick to judge that claims that someone “knows” something are incompatible with other claims (made in completely different conversations) that the same person “doesn’t know” that same thing, we *are* quick to issue such a judgement about many pairs of claims such that it seems that if contextualism were correct, we wouldn’t or shouldn’t be so quick to so judge.

But then I think a good reply to this objection, and perhaps to objections from disagreement and also retraction generally, will be an analogue of my reply to the objection from belief reports (at 2009: 168-170): to argue that the evidence being used against contextualism about “knows” reports “behavior” of “knows” that is like analogous behavior exhibited by terms that contextualism is acknowledged to be true of—and my favorite term to use here is “tall.”

Now this may be a surprising move, because when MacFarlane raises the problems of disagreement and retraction for epistemic contextualism, as at (2014: 181), he does so specifically by *contrasting* how “knows” behaves with the behavior of “tall.” “Here there seems to be a real contrast between ‘know’ and context-sensitive words like ‘tall’,” he writes. But here’s the thing: when he goes on to make his comparison, he doesn’t do it as I compare the terms on how they function in belief reports—and then as I go on in subsequent sections to compare their behavior in speech reports and then in connection with claims (closely related to issues of retraction) like “I never said that!” What I do when I make the comparisons, trying to line up apples with apples and oranges with oranges, is to consider how the terms are used in the various settings in cases where, under contextual analysis, there are “clear and marked differences” in content involved, verses

cases where there isn't such a clear and marked difference, but, on contextualist analysis, it seems one should suspect there at least might well be subtle differences in content.

What may be going on here where the critics think that “tall” clearly behaves quite differently from “knows” is that they think there aren't, in the case of “tall,” such potential subtle differences to worry about. Here, they display the attitude I complain about in (DeRose 2008), where it's assumed that when speakers say that someone is “tall,” they can completely specify what they meant by that by the use of a handy little “for an F” clarification, which specifies the comparison class against which their use of “tall” was issued, and which they clearly had in mind when they made their claim. But even where speakers do have the same F in mind, specifying the same comparison class doesn't settle it that the two speakers were using “tall” with the same content (and thus were saying incompatible things if one was saying the person being described is “tall” and the other saying that person is “not tall”). For the likes of “tall for a fifth grader” are themselves context-dependent. As John Hawthorne observes:

Note that, even having a fixed reference class, there may be context-dependence generated by the need for a standard of application—some kind of threshold on a scale—relative to that reference class. (To be tall do you have to be taller than 80% of the reference class? 85%?) (2004: 53 – 4)

After listing several other Hawthorne-like options, my take on this issue:

Many of our uses of 'tall' seem too vague to be susceptible to the numerical precision displayed by some of the above options; we often have nothing so precise in mind. But our at least apparently more vague uses of 'tall' vaguely approximate different of these options on different occasions. (2008: 146)

And that context-sensitivity in the content of “tall for an F” gets inherited by simple uses of “tall.” What did you mean by “Sally is tall”? Tall for a fifth-grader? But how tall for a fifth-grader did you mean? That can be different on different occasions.

So let’s say an ordinary speaker of English hears one other speaker describe Sally as “tall,” and another speaker, in a completely different conversation, says that Sally is “not tall,” where there is no clear, obvious difference in content in their uses of “tall”—perhaps there’s reason to think both speakers had the same comparison class in mind. Here I think ordinary speakers are quick to judge that the two speakers are in disagreement. But now, we seem to have grounds, as we have for the analogous case with “knows,” for worries that we shouldn’t be that quick to judge the claims to be incompatible if contextualism is true. For shouldn’t our judge be worried about the chances that the two speakers, who are after all engaged in completely different conversations, are setting the threshold for what counts as “tall for a fifth-grader,” and thus for what, on their current uses, counts as “tall,” at slightly different places? And if the denier of Sally’s “tallness” is setting the threshold a bit higher than is the affirmer, the two claims really aren’t incompatible with one another. The answers to this worry—like that maybe there’s some default value for this threshold, so, in the absence of something to move us off the default, one *can* safely assume the speakers’ contents line up; or like that maybe it’s somehow perfectly fine to judge that two claims are incompatible just so long as it safe enough to suppose that they’re quite close to being incompatible, as will be the case where whatever differences there are small and subtle—seem about as available for judgments about the incompatibility of claims involving “knows” as in the case of claims involving “tall.”

So, what about cases where, on contextualism analysis, there *is* a clear difference in content? Well, this a place where we’re likely to have disagreement over the phenomena of

disagreement, but *I* have always taken it that, even in the case of “know,” it *won't* be so clear that the two claims are incompatible with one another. So, for instance suppose (as at 2009: 4-6) that Thelma and Louise were both at the office together earlier in the day, and though neither of them saw John there, they together got excellent reasons to suppose that he was in: they both saw his hat hanging on his office door, which has always been a very reliable indicator of John's presence, and they together heard the very reliable Frank report that John was there. Now, Thelma's at the tavern, collecting on a \$2 bet she made that the often absent John would indeed be at the office today. She's asked how she knows John was in, and when she cites her grounds, they're immediately accepted, and the winnings are paid. Then Thelma is asked if Louise, who also bet that John would be in, knows that John was at the office, so she'll know to collect her money. Thelma replies, “Yes, Louise knows, too.” Meanwhile, Louise is talking with the police, who are investigating John's whereabouts earlier in the day. They have made it clear that this concerns some horrible crime, and they are looking for conclusive grounds for placing John at the office earlier in the day. When the police ask whether she knows that John was at the office, Louise says, “Well, someone said that he was in, and I did see his hat hanging on his door, and I can testify to what I heard and saw, but, no, I don't that John was there.”

So, here Thelma says that Louise “knows” that John was at the office, while Louise herself says that she doesn't “know” this. Is it clear that these two claims are incompatible with one another? I say no—and I'm just a bit tempted to say that it's a *somewhat* clear no. But, you know, I am a contextualist, so maybe I'm unusually inclined to that negative answer? But what would ordinary speakers, with no professional stake in the matter, say? In what follows in this paragraph, I will slightly adapt what I write at (2009: 158-159) to make the case that it's at least far from clear that ordinary speakers will judge there is a contradiction: If you teach a class (perhaps in

epistemology or philosophy of language) in which the issue is relevant, you can try the issue out on your students: Present this case to them, and see whether they think the two claims are incompatible (so that they can't both be true). When seeking intuitions, it's often important that issues be presented as fairly and as neutrally as possible, because, as we all know, the manner of presentation can greatly affect the responses one receives. But I think you can forget all that in this case. Present the issue in a way that creates a strong bias against the contextualist answer. Indeed, rather than asking in any fair way whether there's incompatibility here, you can authoritatively assert that there is, and then ask for feedback on a related question. Try this, for instance: "Thelma says that Louise 'knows' that John was at the office; Louise says she doesn't 'know' that. *Obviously*, these assertions are incompatible with one another, so they can't both be true. So, I ask you: which speaker is making a true claim, Thelma or Louise? If my past experience is any guide at all, in almost every class of sufficient size, you will find that some student will overcome the strong anti-contextualist bias of the presentation, insist that both speakers are asserting truths, and then, if you're at all lucky, will even present some proto-contextualist analysis of the situation to explain how this can be. And once this happens, a significant proportion of this brave student's classmates will join her in her verdict that both are asserting truths (though some squabbles may well break out within this camp over the details of the proto-contextualist analysis they're using to explain the situation). In fact, it's a decent bet that significantly more students will follow this contextualist suggestion than will stick with you on the issue of incompatibility. (Though maybe there's something about me personally which makes my students delight in disagreeing with what they take to be my positions?) But for current defensive purposes, it's enough that there is no clear intuitive winner here. And that, at least, seems quite safe.

Well, I have been going on for too long about disagreement—though I suppose some length is to be expected, given that my theme is that you really have to look at various relevant cases carefully. I do think that a careful look at retraction would yield a very similar outcome, but on retraction, I will cut to the chase, and let’s just look at a case involving some skeptic, but not a global, philosophical skeptic, with some utterly fantastic hypothesis, like that I’m a BIV, but a skeptic who introduces the more down-to-earth possibility that the animals I took to be zebras in the zoo were cleverly painted mules. And let me dig up my response to the supposed retraction data from way back in (1992),⁸ to show that this is how I’ve always felt about whether we should retract earlier positive claims to know after a skeptic has clearly (at least under contextual analysis) raised the epistemic standards on us. I’m responding to earlier work by Palle Yourgrau, citing (but not by this name) supposed retraction data back in (1983):

Yourgrau writes of the zebra case, “Typically, when someone poses a question regarding whether we really know that P obtains rather than some alternative to P, if we cannot satisfactorily answer the question, we conclude that our earlier claim to know was faulty” (p. 183). But do we? We do not stubbornly repeat ourselves, to be sure: “Still, I know that it is a zebra!” We *might* even say, “I don’t know” or “I didn’t know.” All of this the contextualist can handle. But do we (or should we) admit that our earlier claim was *false*? I am on the witness stand being questioned. . .

I then imagine first claiming to “know” there were zebras at the zoo, then deciding to admit, under questioning, that I didn’t know that they weren’t cleverly painted mules, and so didn’t know they were zebras. I then imagine the lawyer jumping all over me, claiming that I’ve contradicted myself. And, well, I was having none of it. My response: “Surely something is wrong in this dialogue. My

⁸ Though I now see this material was in my dissertation, so it goes back to (1990: 35-37).

lawyer should object” (1992: 96). I certainly above rejected Yourgrau’s claim that under the circumstances in question “we conclude that earlier claim was faulty.” And that’s how it’s always struck me, not just about this case, but quite generally: when skeptics, even somewhat down-to-earth ones, press us to not only admit that we don’t know, but to admit our earlier claim was faulty, they are overstepping—and if this isn’t a case of clearly overstepping, then it’s least far from clear that it isn’t overstepping. And if the speakers that the skeptics are pestering in the relevant situations admit their earlier claims were faulty, then they are overreacting, if not clearly, then at least they’re not clearly not overreacting. Which is to say, I’ve never found the supposed retraction data at all convincing.

Finally, though it doesn’t make her banner statements of what she will do or has done in her comments, I do want to respond on the matter of the basic way we can engage skepticism. First, I want to agree with Brendel’s statement: “To my mind, no epistemic view can definitively show that radical skeptical hypotheses do not hold.” In fact, you can take the “definitively” out of it, and I would still agree: so far as I understand what it is to “show” something, we can’t show this at all. This is tied up with my admission, that I make some effort to show is shared by many other current and recent wrestlers with skepticism, that radical skepticism cannot be “refuted,” as that term is best understood. But then how can skepticism be resisted? It’s here that my notion of “defeating” the skeptic provides a good option. For if, as I argue, the best the skeptic has going for them is an argument from intuitively plausible premises to their skeptical conclusion, together, perhaps, with a damage-controlling explaining away of the intuitions that are hostile to them, then if we can incorporate the intuitions the skeptic utilizes into a non-skeptical solution to the puzzle that their argument confronts us with, and show that ours is the best way to make sense of all the intuitions involved—if, that is, we can “defeat” the skeptic, as I use the term—, we will have

shown that our best, most rational response to the skeptic's argument is not to accept their solution or their conclusion. That kind of "defeat" is what I claim to execute against the "bold" skeptic. Though this use of "defeat" is a technical one, it is a technical use that I think merits its label.

4. Reply to Baumann: Lotteries

As Peter Baumann notes, he and I agree about a lot of things where it comes to the epistemology of lotteries, so we will certainly disappoint those who have come to see a good fight. But answering some of the good questions Baumann raises will provide me the opportunity to explain a few things that I should have, but did not, explain in my chapter on lotteries.

Baumann sets up the lottery puzzle nicely with:

Jackie's Puzzle

- (i) Jackie does not know that (a) she won't win the lottery;
- (ii) Jackie does know that (b) she won't be able to afford a much bigger apartment any time soon;
- (iii) If Jackie does not know (a), then she does not know (b).

And then he wonders why I didn't set up the puzzle so nicely. As I set up the puzzle in Chapter 5, I am dealing with an intuition in something like Baumann's (i), with an intuited lack of knowledge of something like (a), but my puzzle pairs that intuition against the intuition that we do know, not something like Baumann's (b), which is nicely connected with (a), but with an intuition that we do know who won a basketball game based on reading the result in a newspaper. Why choose such a disconnected piece of apparent knowledge in setting up the lottery puzzle?

So, the first thing to say is that Baumann is right that the way he has set up the puzzle is generally better. And I do often deal with it in a way like Baumann sets it up: see my E1/O1 pair of cases at the start of Chapter Six (2017: 177-178). My

O1. I won't be able to repay my loan by the end of the year,

which pairing I've been often using since (DeRose 1996: 573), is in the relevant respects like Baumann's (b).

So why the newspaper-based knowledge in my set-up of the problem in my chapter on the lottery? The short answer is that I'm following Gilbert Harman's set-up of the problem (Harman 1968). But Baumann would realize that, and it only pushes back the question to: well, since it is worse, why follow Harman's set-up? And the answer is that I do so as part of an effort to try to take out a problem with the philosophical discussion of lotteries by its very root, and thereby make room in people's thinking about the puzzle for a solutions on which one does know that one has lost the lottery.

I call the puzzle we are wrestling with here the "Harman lottery puzzle," because it seems to trace back to (Harman 1968). (Due to temporal priority, I think the simple title "the lottery puzzle/paradox" should refer instead to what I call the "Kyburg lottery puzzle" (Kyburg 1961).) Harman set up his puzzle by comparing our apparent lack of knowledge that we've lost the lottery with our apparent knowledge of things that we apparently gain by being told some fact or by reading it in the newspaper (the basketball aspect, with my beloved Bulls beating the evil Knicks, is my embellishment). But Harman didn't just set up the puzzle in that way; he also did much to set the assumptions which treatments of the puzzle seem to work with. Harman took the two judgments mentioned above—that we do know things we read in the newspaper, but not that we've

lost the lottery—to be “natural” and “ordinary” judgments (which characterization I suppose I would agree with) and then (to get to the part I don’t like) seems to set out more to understand them than to question them. Here’s how I characterize what seems to be Harman’s procedure:

Harman takes it as a working assumption (that’s not his own description, but I believe it accurately describes his way of proceeding) that the two “natural” “ordinary judgments” here are correct, to see where that working assumption leads, and if it leads to a sensible enough picture of what’s going on, to take that as reason to accept the picture that emerges. And, as things turn out, Harman thinks this procedure does lead to a sensible destination, where we learn what kinds of grounds underwrite knowledge of flat-out claims (as opposed to claims about what is probably the case) of the likes that the Bulls won or that one has lost the lottery. (2017: 133)

And as I further report:

Harman’s basic approach has been followed in most of the subsequent literature on the epistemology of lotteries, with most epistemologists working on the topic pretty much taking it for granted, or at least taking it as a working assumption, that subjects don’t know that they’ve lost the lottery (in the standard situation), while they do know various ordinary things (and so the denial of knowledge in lotteries isn’t part of some general skepticism). (2017: 133-134)

And Baumann agrees with this assessment of the state of the discussion; he calls these assumptions part of the “current orthodoxy”—and one gets the feeling that “current” is in there not so much because he thinks things were different at an earlier time, but out of a hope that things will change in the future.

For Baumann and I both resist orthodoxy here, and like solutions on which you do know you've lost the lottery—and I accept such a solution. Well, as I like to add: unless you are the winner, in which case you're rational to think you know you've lost. And, well, I also add that you know you've lost *by ordinary, moderate standards for knowledge*, for on my contextualist solution, which is in this way like my solution to the skeptical puzzle, you don't know that you've lost the lottery by the very standards for knowledge that tend to get invoked by saying that you do or that you don't "know" that you've lost the lottery, and I use this to explain why you seem not to have knowledge here. But though I hold the contextualist version of heterodoxy, I feel a kinship with invariantist heretics who hold that you *simply* do know that you've lost the lottery,⁹ to the extent that if I were forced to abandon contextualism and my contextualist solution to the lottery puzzle, I'd go for this invariantist position.

And I've long felt that a big part of what's been working against the positions that one does "know"—either simply or by moderate standards for knowledge—that one has lost a lottery is the long-standing assumption that seems to go back to Harman that this is not the way to go. What's more, I think there are good grounds to be found in (Harman 1968) that he should have set up the puzzle, not as a couple of claims that seem to be in tension with one another (that we do know what we read in the newspaper, but that we don't know that we've lost a lottery), which tension one should seek to resolve so as to hold on to both claims, but as three claims, all of which are quite plausible, and so all of which we should try to do justice to: the two already mentioned, plus the comparative claim that if you don't that you've lost the lottery, then you don't know what's

⁹ Here I use "simply know" as I do in the books and as I explain at (2017: 66, n.1). To "simply know" something, in this sense, is to be such that any speaker of ordinary English, whatever their context, would speak the truth in saying that you "know" the thing in question, and would be saying something false if she denied that you "know." This could be either because invariantism about "know(s)" is correct and one meets the standards that invariably govern knowledge attributions, or because, though contextualism is true, one meets all the allowable standards.

reported in the newspaper. Of course, since this third thing is saying that the first two can't both be right, the way in which we "do justice" to all the claims won't be by straightforwardly endorsing all three of them: we will have to engage in some kind of intuitive "damage control." But I believe Harman's original discussion gives us good grounds for taking the comparative claim to be intuitively quite powerful indeed, even as he has set up the comparison (this would be even stronger if the positive knowledge claim in question were one more connected to the lottery, as Baumann and I agree is the better way to set up the puzzle), and that plus the ways that the intuition that we don't know we've lost the lottery does itself prove to be rather "flighty" in some situations, as I explain in sections 5.16-5.18, make for a very strong case for setting up the puzzle in the way I advocate, which opens up room for the kind of solution that Baumann and I like. That's why I spend so much space on the set-up of the problem, and so much attention to how Harman set it up. But I should have noted that Harman's choice of things that we intuitively do seem to know to compare with our apparent lack of knowledge in the lottery situation is different from what seems to be the better way of setting up the puzzle.

In a note (2017: 154, n. 10), I explain why I am not using the possible worlds formulation of the Rule of Sensitivity in addressing the lottery puzzle. As I explain, in the first paragraph of Section 1.12 of SSP, I present my proposed "Rule of Sensitivity" in the way it appeared in my dissertation (DeRose 1990): Making use of the notion of comparative strength of epistemic position, but without appealing to possible worlds. (In my dissertation, I just used "comparative conditionals" to get a fix on comparative strength of epistemic position.) Then in the second paragraph of that section, I proceed to explain the notion of strength of epistemic position and the Rule of Sensitivity in possible worlds terms. In Chapter 5 of the book, I apply my solution to the AI skeptical puzzle to the lottery puzzle, but in its original form, without the explanation in terms

of possible worlds. The original form transfers nicely, but the later elaboration raises some tough questions when applied to the lottery puzzle, because it involves supposing that the possibility that I have won the lottery is very distant. Baumann wonders whether this is a problem: “Why not consider it [winning the lottery] a close possibility?” The problem is that my solution requires that one’s belief that one hasn’t won in the standard lottery situation to be both insensitive and to be a belief one is in a strong epistemic position with respect to (strong enough to satisfy moderate standards for strength of epistemic position), and those can’t both be met on the possible worlds construal of the Rule and of strength of epistemic position, if the possibility of one’s winning is a close one. For on the possible worlds approach, the insensitivity of my belief that I haven’t won would mean that in the closest worlds in which I have won, I still believe that I haven’t won. But then if the closest worlds in which I have won are very close, then I get the matter of whether I’ve won wrong in very close possible worlds—which would mean I’m not even in a moderately strong epistemic position with respect to the matter. My solution to the skeptical puzzle transfers to the lottery puzzle only if the possibility that I’ve won is like the possibility of my being a BIV in being a quite distant or remote one. It needn’t be *as* distant the fantastically remote BIV possibility, but it does have to be distant enough to be outside of moderate standards for possibilities of error that are too close for “knowledge.” Rather than working through these problems (perhaps in the direction pointed to toward the end of that note), I just gave my solution to the lottery puzzle in the original, non-possible-worlds form of my construals of strength of epistemic position and the Rule of Sensitivity.

One of the strongest objections to solutions on which one claims we do know that we’ve lost the lottery in standard lottery situations, or that we know this by moderate standards for knowledge, is generated by closure principles. For if I know that I’ve lost (or “know” by ordinary

standards that I lost), it seems that I should also similarly know that Abe has lost, and that Ben has lost, and that Claire has lost..., where Abe, Ben, and Claire are all in fact fellow-losers. The problem here is not that we might think that you know of the winner that she has lost: where you judge that you know that someone has lost where they are in fact the winner, we say that there you didn't know, because your belief wasn't true, but you were rational to think you knew (or that you knew by moderate standards). The problem comes if you start putting these pieces of knowledge (or knowledge-by-moderate-standards) together, and start deducing that none of the lottery players in very large groups of players have won, for you seem clearly not to know such things by any reasonable standards for knowledge. Baumann and I agree that the problem here is that the objection is making use of a mistaken form of the closure principle for knowledge, but I do think he mischaracterizes my position on closure principles. Realizing that these will likely not get all the details right, Baumann formulates these two rough principles with the labels given here:

(Closure) If S knows that p , and if S knows that p entails q , then S knows that q

(Multi-Premise Closure) If S knows that p_1 , knows that p_2 , ..., knows that p_n , and if S also knows that p_1, p_2, \dots, p_n entail their conjunction, then S knows the conjunction $p_1, \text{ and } p_2, \dots \text{ and } p_n$ ¹⁰

And Baumann writes:

¹⁰ This principle would probably be better labelled "Multi-Premise Conjunction," and I think Baumann may here have intended to instead list this more general principle as what I argue against:

(Multi-Premise Closure) If S knows that p_1 , knows that p_2 , ..., knows that p_n , and if S also knows that p_1, p_2, \dots, p_n together entail q , then S knows that q

What I go on to write above in the text would apply to this more general formulation of (Multi-Premise Closure) as well.

DeRose argues (see his sections 5.20–5.22), convincingly, that (Multi-Premise Closure) is not plausible and way too strong for an acceptable closure principle. One should reject it and restrict one’s closure principles accordingly to something like (Closure).

Now I do argue that the likes of Baumann’s (Multi-Premise Closure) are too strong (though I tend to aim at the more general principle; see note 10, below), and I do advocate retreating to a principle that does not generate the objection we’re considering to the kind of solution to the lottery puzzle that I accept. But my suggested retreat isn’t to something like Baumann’s (Closure). That makes it sound like the problem with (Multi-Premise Closure) is that it applies to multi-premise deductions. But I think that the real problem here is the problem of the “aggregation of risk” (to use Hawthorne’s term), and that this can affect even single-premise deductions. How can there be a problem of aggregation of risk even where there’s just a single premise? Because one is relying on both the truth of the premise, and its connection to the conclusion, and the risks (or “micro-risks,” as I put it in the book, explained at 2017: 170-173) of error involved here can aggregate, and could result in failing to know the conclusion of the deduction. I think the enemy here isn’t multi-premise forms of closure (though the problem shows up more clearly in deductions from many premises), but “Oxford closure,” as I call it: closure principles that do not include a fix for the problem of the aggregation of risk. I give ideas of what a fix here may look like at (2017: 284, n. 12), and at and around the main text of pp. 283-284 where that note attaches. But at any rate, it’s “Oxford closure” that generates the problem, and that form of closure does seem implausibly strong, as I argue in sections 5.20-5.26 of the book. (And in case this paragraph might lead one to think otherwise, I think Baumann and my views on closure and how it relates to the lottery puzzle are quite close to one another.)

Baumann has very interesting things to say about the connection between knowledge and the closing of inquiry. Since he asks, I'll close with a brief explanation of how I would be inclined to approach this. I am dubious of principles like the ones Baumann calls (CK) and (KC) in roughly the way that, and for the reasons that, in Chapter 7 of my first volume (2009), I'm dubious of relativized principles connecting knowledge with action; see especially section 7.10 of the first volume at (2009: 262-269). In fact, the issue of whether a subject can take proposition as a reason for action and whether the subject can or should close inquiry on that proposition seem closely related, and much of my thinking about these issues' alleged connections with knowledge would transfer from the one case to the other. So, I will think that, when you're discussing what a subject does or does not "know" in connection with the issue of whether they ought to or may close inquiry on a topic, there is pressure, other things being equal, for you to set the epistemic standards at a level appropriate to the subject's practical situation, so that whether they "know" or not, at the standards you have set, well marks out the division between their being in a position to close off inquiry or not. This I think can explain some of the considerations that might tempt one toward principles like (CK) and (KC). But since this is defeasible pressure, those principles, nor relativized versions of them (versions relativized to varying epistemic standards) don't really hold. Since claims like "She knows that p, but she ought to continue to look into whether p" seem generally more awkward than do claims like "Though she doesn't yet know that p, she ought to stop looking into the matter of whether p is the case," I take it that the pressure to set the standards so that the subject's meeting those standards is sufficient for her to close off inquiry is stronger than is the pressure to set them so that meeting the standards set is necessary for closing off inquiry. However, the cases devised by Jessica Brown (2008: 176-177), that I discuss at (2009: 265-269), which involve extremely high stakes, provide cases that seem to me, when adjusted so as to involve

closing off inquiry, to show that it's sometimes not all that bad to talk about subjects as if they "know" but should continue to look into the matter. Consider Brown's SURGEON case, where it is immensely important that the surgeon remove the correct kidney (the patient's diseased kidney, rather than their only good one!) from their already anaesthetized patient. Indeed, in that case, we barely have to change what's said, because what's said as Brown constructed the example concerns whether to inquire further (by checking one last time) as well as (the very closely related matter of) whether the surgeon is ready to act. But here is a very slight adjustment:

Student: I don't understand. Why is she looking at the patient's records? She was in clinic with the patient this morning. Doesn't she even know which kidney it is?

Nurse: Of course she knows which kidney it is. But imagine what it would be like if she removed the wrong kidney. She should check the patient's records one last time before operating.

Here, what the nurse says sounds alright to me—though my discussion of the very nearby thing that Brown has the nurse say, on which she might do even better precisely by avoiding saying that the surgeon "knows," would apply here. So, let me make the slight adjustment. This reply seems available:

Nurse: Imagine what it would be like if she removed the wrong kidney! This is too important for her to be relying on her memory from this morning. She should review the patient's records one last time to be absolutely sure that she is removing the correct kidney before she operates.

To mimic what I said about this case's bearing on the knowledge-action connection: it's interesting that this can seem such a good reply, since it sidesteps the question the student explicitly asked

about whether the surgeon “knows,” and instead combines an implied negative judgement about the surgeon’s epistemic position put in the different terms of her not yet being “absolutely sure” with an evaluation of whether the surgeon should continue inquiry by checking the records one last time. This reply seems to me better than the first one we considered—and precisely because it doesn’t combine a positive evaluation of whether the surgeon “knows” with the judgment that she ought to check again. But since the first reply also seems allowable, I think this all supports the thought that there is no real principle at work here, but, at least on a contextualist approach, some defeasible pressure on how the standards for knowledge should be set.

Finally, and on a very closely related note, while I also like much of what Baumann argues in (the quite friendly) section 6 of his comments, I think the contextualist way that I take a solution on which one does know that one has lost the lottery—on which I only hold that one knows this by moderate standards for knowledge—helps to take the sting out of the pointed question Baumann asks at the start of that section: “Why do newspapers print lottery results if all those with losing tickets already know that they have lost?” For I can answer: “So that those losers can come to know that they have lost by the unusually elevated standards often used in discussion of lotteries.”

References

- Blome-Tillmann, M. 2009. “Contextualism, Safety, and Epistemic Relevance,” *Philosophical Studies* 143 (3): 383–394.
- Brown, J. 2008. “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning,” *Noûs* 42: 167–189.

- DeRose, K. 1989. "Reid's Anti-Sensationalism and His Realism," *Philosophical Review* 98: 313-348.
- DeRose, K. 1990. "Knowledge, Epistemic Possibility, and Scepticism," PhD. dissertation, University of California, Los Angeles. University Microfilms International.
- DeRose, K. 1992. "Contextualism and Knowledge Attributions," *Philosophy and Phenomenological Research* 52: 913–29.
- DeRose, K. 1996. "Knowledge, Assertion, and Lotteries," *Australasian Journal of Philosophy* 74: 568–80.
- DeRose, K. 2004. "Sosa, Safety, Sensitivity, and Skeptical Hypotheses." In J. Greco (ed.), *Ernest Sosa and His Critics*, 22-41. Oxford: Blackwell Publishing.
- DeRose, K. 2005. "Direct Warrant Realism." In A. Dole and A. Chignell, (eds)., *God and the Ethics of Belief: New Essays in Philosophy of Religion*, 150-172. Cambridge: Cambridge University Press.
- DeRose, K. 2008. "Gradable Adjectives: A Defence of Pluralism," *Australasian Journal of Philosophy* 86: 141–60.
- DeRose, K. 2009. *The Case for Contextualism: Knowledge, Skepticism and Context, Volume 1*. Oxford: Oxford University Press.
- DeRose, K. 2017. *The Appearance of Ignorance: Knowledge, Skepticism and Context, Volume 2*. Oxford: Oxford University Press.
- Harman, G. 1968. "Knowledge, Inference, and Explanation," *American Philosophical Quarterly* 5: 164–73.

- Kyburg, H.E. 1961. *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.
- Lewis, D. 1979. "Counterfactual Dependence and Time's Arrow," *Noûs* 13: 455–76.
- MacFarlane, J. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. Oxford: Oxford University Press.
- Manley, D. 2007. "Safety, Content, Apriority, Self-Knowledge," *Journal of Philosophy* 104: 403–23.
- Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Plantinga, A. and Wolterstorff, N (eds.) 1983. *Faith and Rationality: Reason and Belief in God*. Notre Dame: University of Notre Dame Press.
- Richard, M. 2004. "Contextualism and Relativism," *Philosophical Studies* 119: 215–242.
- Sosa, E. 1999. "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13: 141–53.
- Sosa, E. 2004. "Replies." In J. Greco (ed.), *Ernest Sosa and His Critics*, 275-325. Oxford: Blackwell Publishing.
- Vogel, J. 1987. "Tracking, Closure, and Inductive Knowledge." In S. Luper-Foy (ed.), *The Possibility of Knowledge: Nozick and His Critics*, 197-215. Totowa, NJ: Rowman & Littlefield.
- Wolterstorff, N. 1983. "Can Belief in God Be Rational if it Has No Foundations?" In A. Plantinga and N. Wolterstorff (eds.) *Faith and Rationality: Reason and Belief in God*, 135-186. Notre Dame: University of Notre Dame Press.
- Yourgrau, P. 1983. "Knowledge and Relevant Alternatives," *Synthese* 55: 175–90.